ISCR

Institute for Scientific Computing Research

ISCR Subcontract Research Summaries

CASC Center for Applied Scientific Computing

Scalable Nonlinear Iterative Methods for Partial Differential Equations

Xiao-Chuan Cai
University of Colorado, Boulder

Summary:

We conducted a six-month investigation of the design, analysis, and software implementation of a class of singularity-insensitive, scalable, parallel nonlinear iterative methods for the numerical solution of nonlinear partial differential equations. The solutions of nonlinear PDEs are often nonsmooth and have local singularities, such as sharp fronts. Traditional nonlinear iterative methods, such as Newton-like methods, are capable of reducing the global smooth nonlinearities at a nearly quadratic convergence rate but may become very slow once the local singularities appear somewhere in the computational domain. Even with global strategies such as line search or trust region, methods for $F(u)=0$ often stagnate at local minima of $||F||$, especially for problems with unbalanced nonlinearities, because the methods do not have any built-in machinery to deal with the unbalanced nonlinearities. To find the same solution of u^* of $F(u)=0$, we solve, instead, an equivalent nonlinearly preconditioned system $G(F(u^*))=0$, whose nonlinearities are more balanced. In this project, we proposed and studied a parallel nonlinear additive Schwarz-based nonlinear preconditioner and showed numerically that the new method converges well even for some difficult problems, such as high Reynolds number flows, where a traditional inexact Newton method fails by stagnation.

We developed a general theoretical framework of nonlinearly preconditioned inexact Newton methods, which we presented at the CASC-sponsored Workshop for Scalable Nonlinear Solvers in July 2000 and which we have submitted for publication in *SIAM J. Scientific Computing*. As an example of the general framework, we formulated and tested in parallel on the Compass Alpha Cluster a nonlinear additive Schwarz preconditioner, called ASPIN (for Additive Schwarz Preconditioned Inexact Newton). Our test results for a viscous incompressible flow problem showed that the method is fast and insensitive to the high Reynolds numbers, where boundary layers appear. The method has great potential for other nonlinearly difficult problems, such as (1) problems whose solutions have local singularities such as shocks or nonsmooth fronts; and (2) multi-physics problems with drastically different stiffness that require different nonlinear solvers based on a single physics submodel. Although the method offers scalable nonlinear convergence in the limit of "bad parameters" such as Reynolds numbers, as with any Newton method, it needs to be coupled on the inside to a scalable linear solver. There exist many fruitful possibilities for combining ASPIN with the parallel linear solver software being developed at LLNL.

Developing a Tuned Version of ScaLAPACK's Linear Equation Solver

Jack Dongarra
University of Tennessee

Summary:

The LINPACK Benchmark has been used as a yardstick in measuring the performance of the Top 500 installed high-end computers. This benchmark was chosen because it is widely used and performance numbers are available for almost all relevant systems. The LINPACK Benchmark solves a dense system of linear equations. For the Top 500, the benchmark allows the user to scale the size of the problem and to optimize the software in order to achieve the best performance for a given machine. This evaluation does not reflect the overall performance of a given system, as no single number ever can. It does, however, reflect the performance of a dedicated system for solving a dense system of linear equations. Since the problem is very regular, the performance achieved is quite high, and the performance numbers give a good check of peak performance of a system.

By measuring the actual performance for different problem sizes *n*, a user can get not only the maximal achieved performance R_{max} for the problem size N_{max}, but also the problem size $N_{1/2}$ where half of the performance R_{max} is achieved. These numbers, together with the theoretical peak performance R_{peak}, are the numbers given in the Top 500 ranking. In an attempt to obtain uniformity across all computers in performance reporting, the algorithm used in solving the system of equations must confirm to the standard operation count for LU factorization with partial pivoting. In particular, the operation count for the algorithm must be $2/3\ n^3 + O(n^2)$ floating point operations.

As part of this project, we developed a version of the benchmark based on the hardware of the ASCI Blue Pacific system to achieve high performance. It was our goal to produce the fastest implementation that will take advantage of the hardware and software infrastructure on the Blue Pacific machine in achieving this mark. In order to accomplish this goal, our team drew on the expertise in this field developed with our work on the ScaLAPACK and ATLAS projects. We collaborated with researchers at LLNL to produce software for the Linpack benchmark that is fully optimized for the ASCI-Blue Pacific system.

The second part of this project, ATLAS, involves using timings coupled with code generation to automatically adapt linear algebra computations to run optimally on varying architectures. The expertise developed in this project will be leveraged in order to make machine-specific performance tweaks for the computations involved in the benchmark, in particular, tunings for various levels of caches, floating point register and unit usage, and optimal use of threading. We presented a tutorial at LLNL on the ATLAS optimization implementation and described ways to optimize numerical software in general.

Adaptive Mesh and Algorithm Refinement with Discrete Simulation Monte Carlo

Alejandro Garcia
San Jose State University

Summary:

In recent years there has been considerable interest in the physics of mixing a shock-accelerated interface between two gases. This simplified problem is of considerable practical importance for the design of inertial confinement fusion (ICF). Undesirable mixing occurs between the ablator and the fuel as a result of the Richtmyer–Meshkov instability, which occurs when the shock passes through the interface between the two gases. This mixing is detrimental to the efficiency of thermonuclear fusion, potentially reducing the yield below the break-even point.

Due to the inherent nonlinear nature of the problem, numerical simulation is one of the most important tools available to us. However, the broad range of scales, from nanometer (shock thickness) to ICF pellet size, causes a significant number of difficulties both from the algorithmic and the modeling side. The full Richtmyer–Meshkov problem is not only a multiscale problem but is also an example of a multiphysics problem in which different physics is required at different scales (e.g., non-hydrodynamic transport at molecular scales).

This project has developed a new technique for studying the Richtmyer-Meshkov problem, specifically designed for multi-physics: Adaptive Mesh and Algorithm Refinement (AMAR). The AMAR framework extends conventional mesh refinement to molecular scales by introducing a particle algorithm at those scales. Specifically, for gas dynamics, direct simulation Monte Carlo (DSMC) is the preferred simulation technique used for the study of shocks. Being a particle-based algorithm, DSMC is meshless, which is an important advantage to a problem that requires the tracking of the interface between two materials; in continuum simulations the algorithmic requirements for the tracking of the interface introduce limitations on the problems studied, or worse, inaccuracies. DSMC captures all the physics including the effects of molecular diffusion on the interface, the structure of the shock profile, and the dynamics of the shock interface interaction.

Our already demonstrated AMAR methodology for single species gas flows was extended to simulate multiple species. Our DSMC routines have been coupled to an Euler solver designed by the SAMRAI team at LLNL and the adaptive mesh refinement implemented using the SAMRAI framework. The new AMAR code will be validated by comparison with earlier AMAR studies and with DSMC results for the Richtmyer–Meshkov instability obtained jointly with Prof. N. Hadjiconstantinou of MIT.

Research on Parallel Adaptive Finite Element Methods

Michael Holst
University of California, San Diego

Summary:

In this project we studied several fundamental issues arising in the parallel adaptive solution of linear and nonlinear elliptic and parabolic PDEs using multilevel algorithms. We focused our efforts on a new approach described in the paper "A New Paradigm for Parallel Adaptive Mesh Refinement" by Bank and Holst. This new approach requires almost no communication to solve an elliptic equation in parallel, and therefore it has the potential to scale much more efficiently on massively parallel computers than do more traditional algorithms. The algorithm has an inherently multilevel structure, in that a sequence of problems on a refinement hierarchy of meshes is solved during the course of the calculation. In particular, the algorithm has three main components: (1) We solve a small problem on a coarse mesh and use a posteriori error estimate to partition the mesh. (2) Each processor provides the complete coarse mesh and instructs to solve the entire problem, but its adaptive refinement is largely limited to its own assigned mesh partition. (3) A final mesh is computed using the union of the refined partitions provided by each processor. The mesh is regularized into a global conformal mesh, and a final solution is computed using a standard overlapping domain decomposition method or a parallel multigrid method.

In certain circumstances the third step can be avoided, leading to an algorithm with no communication beyond that required to synchronize the processors at the beginning of the calculation. Unfortunately, to avoid this third step and still preserve the approximation quality of the resulting solution, a substantial amount of subdomain overlap must be enforced using heuristic techniques. One of our goals for this project was to study algorithms for performing the third step of the Bank and Holst algorithm efficiently on massively parallel computers with the use of parallel algebraic multilevel and multigraph methods, and to develop techniques for skipping the third step, which would avoid the need to use heuristic overlapping techniques. One breakthrough was the realization that the use of dual-problem error indicators provided an automatic overlapping technique that: (1) involved no heuristics, and (2) is optimal in a certain sense (the amount of required overlap is minimized). Portions of this work were presented at the *Root Finder's Ball* held in Pleasanton, CA, during our visit to LLNL, and portions were also presented in a CASC seminar during the visit.

Automated Diagnosis of Large Scale Parallel Applications

Karen L. Karavanic
Portland State University

Summary:

We are studying the application of experiment management techniques to the problems associated with gathering, storing, and using performance data; our goal is achieving completely automated diagnosis of application and system bottlenecks. In this project we focused on incorporating heterogeneous data from a variety of tools, applications, and platforms, and on designing novel techniques for automated performance diagnosis.

The first task includes representation and API for use and storage of performance data from a variety of tools, applications, and collection methods (tracing or profiling, online or offline, etc.). Eventually we will accommodate data gathered and used in both correctness and performance debugging. We investigated several currently available parallel performance tools: Vampir, Jumpshot, Paradyn, Pablo, AIMS, VT, Scala, and Guideview. We implemented extensions for our own PPerfDB prototype to incorporate resource descriptions and performance data gathered with VT, Vampir, and Jumpshot on the LLNL ASCI Blue and Compass platforms and a locally available cluster of SPARC Solaris workstations. This was the first inclusion of data from trace-based tools into the experiment management prototype.

The second task is developing new techniques for the diagnosis itself. Project member Reena John investigated the Paradyn performance consultant in the context of historical data and heterogeneous resource hierarchies. We began designing a diagnostic framework that will incorporate automatically determined thresholds. This expands upon earlier work in which we developed a new automated diagnosis technique to incorporate prior knowledge.

The Experiment Management paradigm is a useful approach for designing a tool that will automatically diagnose performance problems in large-scale parallel applications. The ability to gather, store, and use performance data gathered over time from different executions and using different collection tools enables more sophisticated approaches to performance diagnosis and to performance evaluation more generally. We will continue our efforts by further development and analysis of online diagnosis using historical data, and by investigating performance data and diagnosis gathered from mixed MPI/OpenMP applications.

Discretization and Iterative Solution Techniques for Elliptic Problems on Non-Matching Grids

Raytcho Lazarov and Joseph Pasciak

Texas A&M University

Summary:

The construction, analysis, and numerical testing of efficient solution techniques for solving elliptic PDEs that allow for parallel implementation are the focus of our research. We have investigated and analyzed a number of discretization and solution methods for solving second-order elliptic problems that include mortar and penalty approximations and domain decomposition methods for finite elements and finite volumes. Techniques for parallel domain decomposition algorithms in the framework of PETSc and HYPRE have been studied and tested. We have implemented hierarchical parallel grid refinement and adaptive solution methods and have tested them on various model problems.

Discretization methods for PDEs on non-matching grids are important in applied scientific computing. They provide great flexibility in the grid generation process, increase the portability of various approximation methods and computer implementations, enhance coarsening strategies in parallel algebraic multigrid methods, and provide a natural and practical means of *a posteriori* error analysis in parallel domain decomposition methods. Our research focuses on the analysis and testing of some methods from this promising approach in order to make competitive implementations in engineering codes for complex applications.

We have completed a comprehensive study and testing of various mortar spaces that can be used in non-conforming domain decomposition methods. This implies that meshes for individual subdomains do not necessarily match along the interface, and the continuity of the solution along the subdomain interfaces is imposed in a weak sense. The mortar space provides the "glue" for this process. However, most of the finite element approximations of the mortar space used in the mortar finite element method have been related to the traces of the finite element spaces on the interfaces, which result in continuous functions. We have constructed mortar spaces by using the dual bases approach and some finite volume ideas. Such constructions are relatively simple, valid for general meshes, and lead to local computations. We have constructed discontinuous local dual basis functions for the mortar space that led to diagonal mass matrices and, hence, simple computational schemes. Our mortar spaces have been tested on a series of three-dimensional problems that include second-order elliptic equations and problems of linear elasticity.

The interior penalty method aims at eliminating the need for additional (Lagrange multiplier or mortar) spaces and imposes (approximately) the required continuity across the interfaces by appropriate penalty terms. In our approach the jumps in the values of the functions along these interfaces is penalized in the variational formulation. For smooth solutions we lose the optimal accuracy due to lower approximation at the interface, but on the other hand we produce symmetric and positive definite discrete problems which have optimal condition numbers.

Summary (continued)

We also addressed the issue of constructing preconditioners for composite non-matching grid discretizations. We proposed and investigated an interface domain decomposition preconditioner, which is spectrally equivalent to the reduced interface algebraic problem. We have tested both the accuracy and the preconditioned method on a series of model problems. A similar procedure for combining mixed finite element problems does not require a mortar space. Methods for solving the corresponding composite problems have been proposed, studied, and tested.

We also worked on the problem of multilevel grid refinement and error control for both finite volume approximations and penalty domain decomposition methods. This work directly connects with research in CASC on developing and testing of parallel algorithms (J. Jones, V. Henson. R. Falgout, U. Meier, C. Tong, and P. Vassilevski). We have developed two- and three-dimensional codes for parallel adaptive grid refinement that produce nested (and matching) grids. The resulting multi-level structure is used to define multigrid preconditioners. In collaboration with P. Vassilevski and C. Tong the software was incorporated into the HYPRE preconditioner library. We use the Finite Element Interface (FEI) specification, which provides a layered abstraction that minimizes the concern for the internal details in the HYPRE library. The resulting software was used for testing various ideas and strategies in the a posteriori error analysis and error control for convection-diffusion-reaction problems in 3D domains with complex structure.

We shall now focus on developing multigrid-like (MG, AMG, multilevel) algorithms that exploit coarse finite element spaces that do not necessarily match across subdomain boundaries. In the first algorithm, which is based on geometric finite element discretization, the refinement, local and adaptive, will be the main direction. In order to allow for efficient parallelization, we will investigate approaches of independent refinement in each subdomain, thus facing the problem of non-matching grids across the subdomain interfaces. In the dual approach, which is related to AMGe (algebraic multigrid method for finite element problems), we need to create coarse spaces (or coarse grids) in parallel, and one possibility is to perform any available sequential coarsening in each subdomain, thus creating non-matching coarse grids across the subdomain boundaries. In either approach we have to resolve the issue of matching the independently created problems at the intermediate grids.

In summary, we are developing efficient parallelizable, scalable algorithms for treating either the geometrically refined spaces or algebraically coarsened grids. The research offers alternatives to the existing efforts in CASC for parallel AMG(e) method, which might be more efficient since it exploits additional properties of the matrices.

The AMR++ Library

Bobby Philip

Front Range Scientific Computations, Inc.

Summary:

Our project focused on the following objectives:

- Development and testing of the AMR++ adaptive mesh refinement library to ensure robustness.
- Development and implementation of multilevel elliptic solvers (FAC, AFAC, AFACx) on adaptively refined Cartesian grids using AMR++.
- Initial extension of these algorithms to overlapping logically rectangular curvilinear grids.
- Proof of the level independence of the condition number of the AFACx algorithm.
- Documentation of the C++ code.

Adaptive mesh refinement (AMR) computations are complicated by their dynamic nature. The development of elliptic solvers for realistic applications is complicated by both the complexity of the AMR and the geometry of the problem domains. The AMR++ library represents an attempt to provide users with the framework to develop their own AMR applications from existing single grid solvers. Ensuring that AMR++ was a robust piece of software, and extending the interface to handle AMR on overlapping grids such as those present within the Overture framework, were principal objectives of this work. To this end, extensive development and testing were performed. The code was also documented in detail. The FAC, AFAC, and AFACx algorithms were implemented on Cartesian grids and the possibility of extending these algorithms to overlapping grids was explored. Initial research in this direction is being currently pursued. Proof that the condition number of the AFACx operator is independent of the number of refinement levels was also developed.

Scalable Solvers and Applications in Advanced Materials

Calvin J. Ribbens

Virginia Polytechnic Institute and State University

Summary:

The purpose of the scalable linear solvers and *hypre* projects is to develop scalable algorithms and software for the solution of large, sparse linear systems of equations using massively parallel computers. In order to meet the needs of a variety of current and future users, an important goal of the *hypre* package and framework is interoperability and extensibility. Toward that end, the PI worked with CASC's Andrew Cleary and Jeff Painter to formalize and extend an object model for *hypre*. From the beginning, *hypre* has been designed in an object-oriented style. The existing library includes several solvers and two user interfaces for defining problem components, all of which reflect a well-defined object model. Hence, the starting point in discussing an object model for *hypre* was the existing one. The goal was to extend this model, making it more precise and consistent, while attempting to reflect the ongoing work of the Equation Solvers Interface (ESI) Forum.

Previously, the *hypre* model was defined only by informal description, convention, and example. The PI worked to formalize and extend the *hypre* object model using the tools and notation defined by the Babel language interoperability project. He also contributed to the initial implementation of several of the classes in this model. A significant advantage of using Babel in the context of *hypre* is that Babel's Scientific Interface Definition Language (SIDL) provides an unambiguous, compilable language for defining an object model, without dictating the implementation language. It proved extremely helpful to have a precise language in which to describe the object model, letting Babel provide (and enforce) the mechanism for implementing the object model in a particular language. As in important side benefit, Babel provides convenient language interoperability, e.g., application codes written in Fortran, C, or C++ can use *hypre* easily, and solver modules written in any of those languages can now be incorporated into the library in a much more straightforward fashion.

While the object *Babel/hypre* model is relatively stable, there is still implementation work to be done, and there are other important issues to resolve. Two of these issues are merging "Babelized" *hypre* with the existing package and responding to the ongoing evolution of the ESI. The *Babel/hypre* work is also being reported in a paper to be presented at the Tenth SIAM Conference on Parallel Processing for Scientific Computing. In this connection, the PI is also organizing a minisymposium at that conference on "Component Architectures for High Performance Scientific Computing."

Summary (continued)

Patrice Turchi is a senior research scientist in the Chemistry and Materials Science (CMS) Directorate at LLNL. He and his co-workers are developing a new, tight-binding-based electronic structure scheme for the prediction of the thermodynamic and physical properties of complex multi-component alloys. The scheme combines molecular dynamics (MD) and monte carlo (MC) simulations so that issues related to different time-scales can be addressed. The energetics which describes the diffusion part of the scheme (MC) is obtained within a real-space electronic structure (ES) method by solving first for the inhomogeneous chemically random alloy in the framework of the coherent potential approximation with a recursion technique, and second for the ordering part of the internal energy with the embedded cluster method and an orbital-peeling technique. Once a new alloy configuration is obtained from MC simulations, the alloy topology is relaxed with MD in the next cycle. The MD-ES-MC loop is iterated until self-consistency is reached and chemical short-range order, energetics, and relaxed topology can be predicted.

Stand-alone sequential codes exist for each of three main components of the MD-ES-MC loops. The MD and ES codes are computationally very demanding, requiring many hours of CPU time for a problem of modest size (about one thousand atoms). During his visit to LLNL, the PI became familiar with these codes, concentrating his efforts on the most expensive step (ES), as implemented in the (sequential) *tbepi code*. This code was first parallelized using OpenMP for the Compaq cluster shared-memory machines at LLNL. Then a distributed-memory MPI version of the code was developed, targeting the IBM ASCI Blue Pacific machine. To take advantage of this architecture a hybrid implementation was developed, with OpenMP-based threads at inner loops and MPI-based problem decomposition at outer loops. A special static load-balancing scheme was developed in order to achieve good performance in the distributed-memory version of *tbepi*. Efficient and scalable performance was demonstrated on up to 2000 atoms on 250 nodes (1000 processors) of ASCI Blue. In addition to the work on *tbepi*, the PI implemented the first MPI-based parallel version of Turchi's MD code *tbmd*.

Development of Object-Oriented Tools for the Numerical Solution of Reactive Flow

Donald W. Schwendeman
Rensselaer Polytechnic Institute

Summary:

Our work developed object-oriented application software, within the Overture framework of codes, for the numerical simulation of high-speed reactive flow. The mathematical model on which the software is based is the reactive Euler equations. The implementation of this model is fairly general and allows for multiple reacting species and reaction rates and a general equation of state with the aim of being able to simulate experimentally observed phenomena in gas or solid explosives. The software is part of the OverBlown package of fluids codes, developed and maintained by Bill Henshaw and the Overture team at CASC. It uses overlapping grids in order to handle general domains and an array class library, developed by Dan Quinlan and the Overture team, which allows parallel processing. The software has been carefully tested for accuracy using existing codes written previously by the author, and it has been used to compute the evolution to detonation of reactive samples subject to various initial conditions and within various confinement geometries. Work on a number of smaller projects has also taken place. These include methods for conservative interpolation on general curvilinear grids, methods for characteristic interpolation, and developing software tools for numerical quadrature on overlapping grids.

A numerical method of solution was implemented for the reactive Euler equations. Within the Overture framework, the numerical method discretizes the equations on a set of overlapping, structured grids that cover the domain of interest. The software package Ogen is used to generate the overlapping grid and provide geometric mapping information for each curvilinear component grid and information concerning the communication of the solution between grids in the overlap region. This information was considered to be given and thus the main task was to implement the numerical method for an individual component grid and then let the existing software, OverBlown, handle the surrounding numerical details (such as interpolation between component grids, the application of boundary conditions, time stepping, graphical interface, and more).

For an individual component grid, an unsplit Godunov-type, finite volume method was used to solve the equations numerically. The approximate Riemann solver is a Roe-type method extended to handle a general equation of state. These ODEs are solved numerically for each grid cell using a Runge-Kutta, variable time step, error control algorithm. This is done to ensure that the (possibly) stiff rate equations are solved sufficiently and accurately.

Summary (continued:

Several example calculations were run to test the software and to study the evolution to detonation of a reactive sample at critical conditions subject to gradients in temperature. The physical motivation for these calculations was to examine the effects of small nonuniformities in the initial state and of the confinement geometry on the severity of the reaction in terms of peak temperatures and pressures. If the gradient parameter for the temperature is too small, then an explosion takes place as a "phase wave" in which a reaction front is born from the initial hot-spot location but travels too fast for significant acoustic information to propagate and strengthen the wave to form a detonation. On the other side, if the gradient parameter is too large, then a local explosion from the initial hot spot does not create a large enough response to initiate significant reaction in the cold neighboring material. In this case, detonation does not occur. In between these extremes, there is a range of values for which the initial gradient leads to a constructive interplay between chemical activity and acoustic propagation resulting in the evolution to a shock-supported detonation. This event leads to high peak temperatures and pressures at the leading edge of the detonation, which are enhanced still further by interactions with a solid confinement

The core elements of the software are now in place and have been tested for accuracy and shown to be useful for numerical simulations. In particular, the software has been used to study the evolution to detonation for a nonuniform reactive sample at critical conditions. While the project has been successful, there are tasks that remain for future work. For example, on the software side, it would be very useful to have an AMR implementation. Work on this aspect of the software is in progress. At present, a Godunov-type numerical flux is implemented with a Roe-type approximate Riemann solver. There are many other choices of numerical flux functions and Riemann solvers that one might use and build into the software as a user option. Additionally, it would be useful to include options for various multiple kinetic reactions, such as chain branching reactions. A long-term goal would be to include diffusive terms in the model equations in order to handle problems involving low-speed flames.

Spatio-Temporal Data Mining of Scientific Trajectory Data

Padhraic Smyth and Scott Gaffney

University of California, Irvine

Summary:

With the increasing availability of massive observational and experimental data sets across a wide variety of scientific disciplines, there is an increasing need to provide scientists with efficient computational tools to explore such data in a systematic manner. For example, techniques such as classification and clustering are now being widely used in astronomy to categorize and organize stellar objects into groups and catalogs, which in turn provide the impetus for scientific hypothesis formation and discovery.

Data-driven exploration of massive spatio-temporal data sets is an area where there is particular need of data mining techniques. Scientists are overwhelmed by the vast quantities of data that simulations, experiments, and observational instruments can produce. Analysis of spatio-temporal data is inherently challenging, yet most current research in data mining is focused on algorithms based on more traditional feature-vector data representations.

Scientists are often not particularly interested in raw grid-level data, but rather in the phenomena and processes which are "driving" the data. In particular, they are often interested in the temporal and spatial evolution of specific "spatially local" structures of interest, e.g., birth-death processes for vortices and interfaces in fluid-flow simulations and experiments, trajectories of extra-tropical cyclones from sea-level pressure data over the Atlantic and Pacific oceans, and sunspot shape and size evolution over time from daily chromospheric images of the Sun. The ability to automatically detect, cluster, and catalog such objects in principle provides an important "data reduction front-end" which can convert 4-d data sets (3 spatial and 1 temporal dimension) on a massive grid to a much more abstract representation of local structures and their evolution. In turn, these higher-level representations provide a general framework and basis for further scientific hypothesis generation and investigation, e.g., investigating correlations between local phenomena (such as storm paths) and global trends (such as temperature changes).

In this work we focused on detecting and clustering *trajectories* of individual objects in massive spatio-temporal data sets. There are two primary technical problems involved. First, the local structures of interest must be detected, characterized, and extracted from the mass of overall data. Second, the evolution (in space and/or time) of these structures needs to be modeled and characterized in a systematic manner if the overall goal of producing a reduced and interpretable description of the data is to be met.

We proposed a general framework for clustering trajectories using probabilistic models of dynamic systems, which allows one to overcome limitations of feature vector-based methods. As such we began looking at various types of dynamic models, for example, autoregressive (AR), moving average (MA), and the more general ARMA model. One of our tasks was to figure out how to simulate direction-focused trajectories from these models in terms of parameter settings. We also looked at learning the parameters of these models given some set of generated data from a known model.

Summary (continued)

An application of trajectory clustering can be found in the clustering of cyclone tracks from meteorological data. We are currently working with a simulated data set for the 1979/1980 winter that gives mean sea-level pressure measurements. We focus on using a bicubic interpolation inside of an iterative scheme to find our minima using a simple gradient descent. First we scan all of the time-sliced images and find all the local minima using a simple sliding neighborhood method. That is, we declare a "pixel" to be at a local minimum if its value is less than all eight of its neighbors. Then we use a simple gradient descent with bicubic interpolation to descend to the point "inside" of the pixel that is at an approximate minimum. This point then gives us our approximate offgrid center of a candidate cyclone. Using the above technique, we processed the data to force all of the grid-based minima to lie in continuous space. We then fed this data into our new tracking software.

We have finished the process of porting our previously developed MATLAB software to our current C++ PC-based platform and, in addition, we have completed the necessary modifications to allow tracking to be carried out using offgrid coordinates. At this time, much of the basic software development has been finished. Much of the future development will be focused on dynamic modelling implementation. We will investigate Kalman filter models.

From here we would like to investigate ways in which the tracking can be integrated into the clustering framework. That is, if we know which cluster an individual belongs to with some probability, then we should be able to more accurately track its future movements. In other words, we believe that instead of two different problems tracking and clustering what we have here is one compound problem that can be solved in an integrated manner.

Algebraic Coarsening Methods for Linear and Nonlinear PDE and Systems

Irad Yavneh
University of California, Los Angeles

Summary:

Algebraic Multigrid (AMG) methods for PDEs were introduced and established in the 1980s. In 1999 Brandt presented a very general approach to algebraic coarsening (including the choice of coarse-grid variables, coarse-grid equations, and inter-grid transfers), which led to an optimal multiscale algorithm for Monte Carlo simulations. The key observation behind this approach, expressed in terms relevant to deterministic PDEs, is the following: Suppose that an appropriate discretization for the fine-grid equations, and some appropriate relaxation method (smoother), have been chosen. Suppose also that an appropriate definition of the coarse-grid variables in terms of fine-grid variables has been selected (e.g., the coarse-grid variables may be a subset of the fine-grid variables, as in classical AMG, or a local average). Define now a compatible relaxation, as an implementation of the above-mentioned relaxation method for the fine-grid variables, modified so as to leave the coarse-grid variables invariant. (For example, in classical AMG this would simply require skipping the variables that coincide with coarse-grid variables.) Then *a general measure of the quality of the set of coarse-grid variable is the convergence rate of the compatible relaxation*. Moreover, fast convergence of the compatible relaxation implies that the dependence of each coarse-grid variable on any other coarse-grid variable decays exponentially (or faster) with the distance between the two, where the distance is measured by coarse-grid mesh-intervals.

The former observation provides us with a tool for gauging the adequacy of our chosen set of coarse-grid variables (and suggests a method for obtaining a set, which is tailored to our needs). The latter implies that a highly accurate coarse-grid approximation to the fine-grid equations can be obtained, where only dependence on a fairly close neighborhood must be explicitly taken into account. It also suggests how such a coarse-grid discretization may he derived by means of solving a local optimization problem. We developed a new compatible relaxation prolongation that seems comparable to classical AMG in its cost of implementation. The new approach offers the potential of generalization to problems where the coarse-grid variables are not (or cannot be) defined as subsets of the fine-grid variables.

As another part of this research, we developed a directly derived coarse grid. To determine the coarse-grid equation for coarse-grid variable, first choose some local fine-grid subdomain, F. We then take a linear combination of the fine-grid equations corresponding to variables belonging to F. This linear combination is itself obviously an exact fine-grid equation. The coarse-grid equation is obtained by ignoring all coefficients corresponding to F-variables. The idea is to choose the coefficients of the linear combination such that the neglected coefficients will be minimal (in some norm), and that the coefficient of the coarse-grid variable will be relatively large so that there will remain an unambiguous correspondence between equations and variables on the coarse grid. The coefficients of the linear combination define a restriction operator. The neglected coefficients of the resulting stencil represent the discrete relative truncation error. We minimize the weighted L_2 norm of the latter, subject to certain constraints. For example, we normally impose that the sum of the neglected coefficients equals zero, which ensures that the coarse-grid approximation of constant functions is exact (since the relative truncation error vanishes in this case).

Summary (continued)

Alternative constraints are also possible, for example. Imposing that all the neglected coefficients be of the same sign, resulting in a coarse-grid operator which satisfies some fine-grid maximum principle. This can be used to obtain monotonic convergence. We implemented this and found the convergence rate was satisfactory, though this approach is expensive. It remains to be seen whether it is advantageous in more general problems.

A central decision in AMG for systems of PDEs is how to choose which equations should participate in the linear combination that makes up each equation in the transformed system. The straightforward approach of including in the linear combination for a coarsened equation all of the fine-grid equations in which it is a dominant variable is unsatisfactory (except for simple problems such as the Cauchy-Reimann equations). We therefore examined a "grey box" version, where we attach a "color" to each type of variable on the fine grid. Then, in the definition of the coarse grid, we apply the usual AMG coarsening process to each color separately (ignoring dependencies on variables of different colors for this process). This approach seems to hold some promise. Indeed, for the Stokes equations with staggered discretization and periodic boundary conditions, this approach can produce an exact decoupling of the equations and the resulting system is solved essentially as fast as the Poisson problem. However, much work is yet needed before a reasonably general "grey box" solver for elliptic systems is achieved. We are currently testing and refining this approach for the Stokes system with boundaries and for linearized Navier-Stokes.

In summary, some new algebraic coarsening methods have been studied, based on ideas related to the condition of fast convergence of compatible relaxation. A generalized AMG approach, where the prolongation operator is derived using compatible relaxation, has been formulated and tested with encouraging results. New methods of direct derivation (i.e., non-Galerkin) have also been introduced and preliminary tests have been made, including for systems of PDEs. Such methods may not be computationally competitive for relatively simple problems, but they can be made quite general.